stochastic games

Review of expected value from basic probability:

Given some numerical outcomes and probabilities for those outcomes, the expected value is the sum of prob x outcome

e.g. • one dice:

	1	2	3	4	5	6
prob	1/6	1/6	-	-	-	1/6

$$\text{exp. val.} = \left(\tfrac{1}{6}\times 1\right) + \left(\tfrac{1}{6}\times 2\right) + \ldots + \left(\tfrac{1}{6}\times 6\right) = 3.5$$

• two dice:

2	3	4	....	11	12
1/36	2/36	3/36	...	2/36	1/36

$$\text{exp. val.} = \left(\tfrac{1}{36}\times 2\right) + \left(\tfrac{2}{36}\times 3\right) + \ldots \left(\tfrac{1}{36}\times 12\right) = 7$$

Exercise: 2 biased coins. Chance of heads is 0.6 for each. Both are flipped, you win $3 for each heads showing. What are your expected winnings?

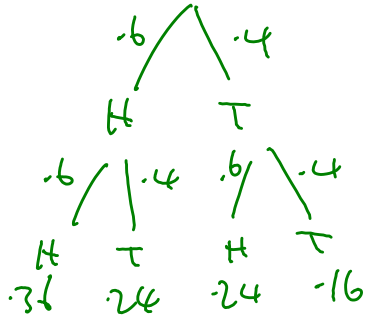

num heads	0	1	2
prob	.16	.48	.36
outcome	0	3	6

exp. val. = $0\times .16 + 3\times .48 + 6\times .36$

= 3.6

Solving a stochastic game:

(1) Use chance nodes (as well as MAX, MIN)

(2) Use minimax, but compute expected value of minimax at chance nodes

e.g. Russell slide 26, Try it with biased coin. and see formal algorithm on p148.

(3) Can cut off and evaluate if desired.

Notes:
- complexity is horrible: $O(b^m n^m)$ (n = number of distinct chance outcomes)
- pruning like α-β is possible but tricky — see homework exercise 5.16.
- can instead evaluate using Monte Carlo simulation...

Monte Carlo simulation

Basic idea: at chance node, do not evaluate all possibilities. Instead, pick a random sample of the outcome, compute minimax for each, and return the average. See p180 and p184 for a few more details.

Final note: The ideas in this lecture contain some promising material for the final project of the course.